

REMARKS

Claims 28-41 are the claim currently pending in the present Application.

Rejection of Claims 29 and 35-41 under 35 U.S.C § 112

Claims 29 and 35-41 are rejected under 35 U.S.C. § 112, first paragraph, as based on a disclosure which is not enabling. This rejection is traversed.

The Examiner alleges that several terms recited by the claims are not sufficiently enabled by the Specification. Applicant herein below cites specific passages of Applicant's disclosure that shed light on the claim terms cited by the Examiner. However it will be understood that the cited passages are provided only by way of example and do not necessarily provide comprehensive definitions of the terms or exhaust the range of possible implementations that the claim terms may support.

In Claim 35

(1) The term "interval hash value computing module" is alleged to be unsupported.

Page 3 of Applicant's disclosure provides that:

In order to rapidly identify the records to be transferred, the databases each support a special facility for efficient computation of a digest (hash) of any specified range of keys, which amounts to a summary of the range. This special facility corresponds to the invention is *summarizable database* abstraction.

(Italics in the original.) Thus, the interval hash value computing module computes a hash or value that serves as a "summary" of the interval, the range of data contained in the database.

(2) The term "a synchronization module coupled to ... the interval hash value computing module" is alleged to be unsupported.

Applicant's disclosure starting on page 17 describes the OSYNCH module, which according to an aspect of Applicant's claimed invention, implements range synchronization for the BxTree database embodiment of Applicant's claimed invention. For example, page 18 provides as follows:

The synchronization algorithm starts by asking both databases to compute a single summary of all records lying in the given key interval. The Get_Interval_Hashes function is invoked for this. The remote summary is transferred to the local side and compared with local summary. If the summaries match, it is concluded that the databases are already synchronized restricted to the giving key interval... These remote summaries are then compared against corresponding local summaries and the operation is invoked recursively for sub-interval [sic] whose summaries do not match.

Further, the Examiner's attention is drawn to Figure 2, which shows the Synchronization Facility 3.

In Claim 29:

The terms "a first" and "a second sub-sub-interval," and "a first" and "a second sub-sub-interval summary hash" are alleged to be unsupported.

As explained in Applicant's disclosure, page 14:

This synchronization operation works in a number of communication rounds. In each round, the key range of interest is partitioned into smaller sub-ranges. For each sub-range, the two databases compute the summary of records lying in that sub-range and one of the databases sends its summaries to the other side. The corresponding summaries from the two sides are compared and the operation is recursively applied to sub-ranges whose summaries do not match. Only those records are transferred from one side to the other which (1) are missing on the other side, or (2) have a mismatching record on the other side. Thus, unnecessary transfer of large amounts of data is prevented.

Thus, the synchronization operation is performed by comparing or by matching a hash value obtained for an interval of data of a remote database with a hash value for a corresponding interval of a second database. If the hash values match, then the data contained in the interval and the corresponding interval do not need to be synchronized. However, if the hashes do not match, then a hash of a sub-interval of the remote database is compared with a hash of a corresponding sub-interval of the second database. This is what is meant by the operation being recursively applied to sub-ranges if summaries or hashes do not match. If these match, then a second sub-interval of the same interval of the remote database is compared with a hash of a corresponding second sub-interval of the second database. However, if the hash of the sub-interval and the corresponding sub-interval do not match, then a hash of the sub-sub interval and a corresponding sub-sub-interval are compared. This recursive process is repeated until the smallest range (which is separately accessible, for example, a data block representing a set of records, a record, or a field of a record) for which the hashes do not match is reached (or until some other desired size range of data is reached).

It is respectfully submitted that a person of ordinary skill in the art would have readily understood each of the foregoing terms based on Applicant's disclosure, and therefore this rejection should now be withdrawn.

Rejection of Claims 28-30, 35 and 37 under 35 U.S.C. § 103

Claims 28-30, 35 and 37 are rejected under 35 U.S.C. § 103 as being obvious from Boothby, U.S. Patent No. 6,141,664 in view of Lomet et al., U.S. Patent No. 5,440,732. This rejection is traversed.

Among the problems recognized and solved by Applicant's claimed invention is that of having to compare most of the data of two databases to synchronize the databases, or having to rewrite entire files specifically provided, in order to synchronize the databases. According to an aspect of Applicant's claimed invention, hash values that "summarize" an interval of the first database and the second database are obtained and these hash values are compared or matched to determine whether a synchronization of that interval is necessary. If the hash values for the interval do not match, then as discussed above with respect to the synchronization facility, sub-interval hash values may be obtained and the process repeated recursively.

For at least the following reasons, Applicant's claimed invention is neither anticipated by nor obvious from the cited references. By way of example, independent claim 28 requires synchronizing based on whether the single summary hash for the records of the interval match the local summary hash of the interval. Further, independent claim 35 requires identifying database records that need synchronizing by comparing the summary hash for the interval with a remote summary hash for the interval.

Boothby is directed to synchronization of incompatible databases based on a date range (Boothby, Abstract); such that incompatible databases (databases with different types of data structures, Boothby, column 1, lines 39-50) are synchronized without overloading the storage capacity of a database contained in a smaller device (Boothby, column 2, lines 1-10). Boothby discloses that the synchronization of the two databases is performed for the current date range and a prior date range (Boothby, column 2, lines 13-23); such that the user specifies the date range for which synchronization is desired when an "incremental" date range synchronization is performed (Boothby, column 5, lines 42-55); or a synchronization from scratch for the entire date range may be performed (Boothby, column 5, lines 5-20).

Boothby does not disclose or suggest obtaining a summary hash value for a record or set of data for an interval of a database, as *inter alia*, required by independent claims 28 and 35. Nor does Boothby disclose or suggest synchronizing databases based on a comparison or matching of such summary hash values, as further required by independent claims 28 and 35.

Lomet discloses a locking strategy for a database employing an index tree (Lomet, Abstract; column 1, lines 6-8) for effective locking and unlocking of a database resource while an update is in process (Lomet, column 3, line 35 - column 4, line 47). Lomet discloses a database indexing approach using a B-tree organization indexing approach, and Lomet discloses a search mode using a hash function, in which a hash value calculated for a search key is used to compute the address of the “bucket” in which the record resides that contains the searched for key value (Lomet, column 40, lines 40-63). Thus, Lomet describes that a hashing algorithm applied to the search key value yields an address of the bucket in which the desired record resides.

Lomet does not disclose or suggest synchronizing database records based on whether the summary hash for the record of the interval matches the local summary hash of the interval, as *inter alia* required by independent claim 28. Further, Lomet does not disclose or suggest identifying database records that need synchronizing by comparing the summary hash for the interval with a remote summary hash for the interval, as *inter alia* required by independent claim 35.

First, Lomet does not disclose or suggest obtaining a summary hash value for a set of data of an interval of a database. As discussed, the hashing function is applied to a search key and yields an address for a record. However, Lomet does not disclose or suggest a hash value obtained for a set of data or data record(s).

Moreover, since Lomet does not disclose or suggest this feature, Lomet is incapable of disclosing or suggesting synchronizing databases based on a comparison or matching of such summary hash values, as *inter alia*, required by independent claims 28 and 35. Therefore, Boothby and Lomet, even if combined and taken together as a whole, do not disclose or suggest the recitation of independent claims 28 and 35.

More generally, Boothby and Lomet belong to the conventional art recognized by Applicant's claimed invention, because Boothby and Lomet do not disclose or suggest the above-discussed problems recognized and solved by Applicant's claimed invention. For instance, the problem of having to compare most of the data of two databases to synchronize the databases, or having to rewrite entire files specified in advance in order to synchronize the databases, is not disclosed or suggested by Boothby and Lomet. Also, at least one of the solutions provided by Applicant's claimed invention, that of synchronizing database data based on hash value summaries for data contained in corresponding intervals of the databases, is not disclosed. Accordingly, Boothby and Lomet do not even remotely disclose or suggest Applicant's claimed invention. Therefore, this rejection should now be withdrawn.

Rejection of Claims 31, 32, 36, 38 and 39 under 35 U.S.C. § 103

Claims 31, 32, 36, 38 and 39 are rejected under 35 U.S.C. § 103 as being obvious from Boothby and Lomet in view of Bendert et al., U.S. Patent No. 5,668,958. This rejection is traversed.

Bdert does not cure the above-discussed deficiencies of Boothby and Lomet as they relate to independent claims 28 and 35, and the Examiner does not allege that it does so. Therefore, since claims 31 and 32 depend from independent claim 28, and claims 36, 38 and 39

depend from independent claim 35, claims 31, 32, 36, 38 and 39 are patentably distinguishable over the prior art for at least the reasons that independent claims 28 and 35 are patentably distinguishable over the prior art. Therefore, this rejection should now be withdrawn.

Rejection of Claims 33 and 40 under 35 U.S.C. § 103

Claims 33 and 40 are rejected under 35 U.S.C. § 103 as being obvious from Boothby, Lomet and Bendert in view of Bozman, U.S. Patent No. 5,089,952. This rejection is traversed.

Bozman does not cure the above-discussed deficiencies of Boothby and Lomet as they relate to independent claims 28 and 35, and the Examiner does not allege that it does so. Therefore, since claims 33 and 40 depend from independent claims 28 and 35, respectively, claims 33 and 40 are patentably distinguishable over the prior art for at least the reasons that independent claims 28 and 35 are patentably distinguishable over the prior art. Therefore, this rejection should now be withdrawn.


Rejection of Claims 34 and 41 under 35 U.S.C. § 103

Claims 34 and 41 are rejected under 35 U.S.C. § 103 as being obvious from Boothby, Lomet and Bendert in view of Hecht, U.S. Patent No. 5,778,375. This rejection is traversed.

Hecht does not cure the above-discussed deficiencies of Boothby and Lomet as they relate to independent claims 28 and 35, and the Examiner does not allege that it does so. Therefore, since claims 34 and 41 depend from independent claims 28 and 35, respectively, claims 34 and 41 are patentably distinguishable over the prior art for at least the reasons that independent claims 28 and 35 are patentably distinguishable over the prior art. Therefore, this rejection should now be withdrawn.

For at least the reasons set forth in the foregoing discussion, Applicant believes that the Application is now allowable, and respectfully requests that the Examiner reconsider the rejections and allow the Application. Should the Examiner have any questions regarding this Amendment, or regarding the Application generally, the Examiner is invited to telephone the undersigned attorney.

Respectfully submitted,


George Brieger
Registration No. 52,652

Scully, Scott, Murphy & Presser
400 Garden City Plaza, Suite 300
Garden City, New York 11530
(516) 742-4343, Ext. 509

GB:ar